Light spanners for Snowflake Metrics

Lee-Ad Gottlieb * Shay Solomon †

arXiv:1401.5014v1 [cs.CG] 20 Jan 2014

Abstract

A classic result in the study of spanners is the existence of light low-stretch spanners for Euclidean spaces. These spanners have arbitrary low stretch, and weight only a constant factor greater than that of the minimum spanning tree of the points (with dependence on the stretch and Euclidean dimension). A central open problem in this field asks whether other spaces admit low weight spanners as well – for example metric space with low intrinsic dimension – yet only a handful of results of this type are known.

In this paper, we consider snowflake metric spaces of low intrinsic dimension. The α-snowflake of a metric (X, δ) is the metric (X, δ^α), for $0 < \alpha < 1$. By utilizing an approach completely different than those used for Euclidean spaces, we demonstrate that snowflake metrics admit light spanners. Further, we show that the spanner is of diameter $O(\log n)$, a result not possible for Euclidean spaces. As an immediate corollary to our spanner, we obtain dramatic improvements in algorithms for the traveling salesman problem in this setting, achieving a polynomial-time approximation scheme with near-linear runtime. Along the way, we also show that all ℓ_p spaces admit light spanners, a result of interest in its own right.

*Department of Computer Science and Mathematics, Ariel University, Ariel, Israel. E-mail: leead@ariel.ac.il.

†Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot 76100, Israel. E-mail: shay.solomon@weizmann.ac.il. This work is supported by the Koshland Center for basic Research.

1 Introduction

Given a complete graph G, a $(1+\epsilon)$-*spanner* for G is a subgraph $H \subset G$ which preserves all pairwise distances in G to within a factor of $1+\epsilon$. Low-stretch spanners with additional favorable properties – such as low degree, weight or small hop diameter – have been the object of much study, in settings such as Euclidean space, planar graph metrics, and metrics with low intrinsic dimension [AS97, AMS94, ADM$^+$95, ACC$^+$96, CG06, GR08a, GR08b, DES08, Sol11, ES13].

A spanner H is said to be *light* if its weight is proportional to the weight of the minimum spanning tree of G, $w(\mathrm{MST}(G))$, and its *lightness* is the constant (or term) multiplying $w(\mathrm{MST}(G))$. A major result of the nineties is that d-dimensional Euclidean spaces admit light $(1+\epsilon)$-spanners, with lightness $\epsilon^{-O(d)}$ [DHN93]. An important result in its own right, the light Euclidean spanner is also a central component in the fastest polynomial time approximation scheme (PTAS) for the Euclidean traveling salesman problem (TSP): Using a light spanner, a $(1+\epsilon)$-approximate tour can be computed in time $2^{\epsilon^{-O(d)}}n + 2^{O(d)}n\log n$ [Aro98, RS98].

The existence proof for light Euclidean spanners is complex. At its core, it relies on the *leapfrog property* specific to Euclidean space. It seems difficult to extend this proof to other natural spaces, and in fact light spanners are known for only a handful of settings. These include planar graphs [ADD$^+$93], unit disk graphs [KPX08], and graphs of bounded pathwidth [GH12] and bounded genus [DHM10]. In fact, a central conjecture in this area asks whether all metric spaces M with low intrinsic dimension admit light $(1+\epsilon)$-spanner. The best lightness bound known for spanners in these metrics is $\Omega(\log n)$ [Smi09, ES13].

In this paper we take a step towards this conjecture, by showing that light spanners exist for snowflake metrics of low doubling dimension; the α-snowflake of a metric (X,δ) is the metric (X,δ^α), with $0<\alpha<1$. Snowflake metrics have been a focus of study in the recent literature [Ass83, GKL03, LMN04, ABN08, GK11, BRS11, NN12, Nei13]. We will give two separate proofs for the existence of light spanners for snowflake metrics. As an immediate corollary, we derive a fast approximation algorithm for TSP on snowflake metrics, which yields a dramatic improvement on what was previously known [BGK12].

The first proof, presented in Section 2, is based on a new observation: For any constant-dimensional vector space V, the edges of a light $(1+\epsilon)$-spanner for space (V,ℓ_2) also form a light $(1+\epsilon^{O(1)})$-spanner for any space (V,ℓ_p), when $p \geq 1$. This result is of independent interest, and also implies an efficient approximation algorithm for traveling salesman in ℓ_p (see Corollary 4.3). Further, combining this result with the fact that the α-snowflake of any metric M embeds with $1+\epsilon$ distortion into $\epsilon^{-O(\mathrm{ddim}(M))}$-dimensional ℓ_∞ [HPM06], we conclude in Theorem 2.5 that every snowflake metric M admits $(1+\epsilon)$-spanners of lightness $2^{\epsilon^{-O(\mathrm{ddim}(M))}}$.

Unsatisfied with the lightness bound provided by Theorem 2.5, we present a second proof in Section 3. This proof is more involved, but yields an exponentially better lightness bound. The proof considers the standard net-tree spanner (NTS), and demonstrates that for snowflake metrics, the NTS is light. In Theorem 3.7, we show that the NTS on a snowflake metric M has stretch $(1+\epsilon)$ with lightness $\epsilon^{-O(\mathrm{ddim}(M))}$. The advantage of this spanner over that of Section 2 is threefold: (1) It bypasses the heavy machinery of the leapfrog property and the use of low-distortion embeddings. (2) The dependencies on ϵ and the doubling dimension in the lightness bound, as well as the leading constant therein, are significantly smaller. This is of particular interest to practitioners in the field. (3) The NTS is the central tool in several spanner constructions that combine small weight with other favorable properties, and the weight bound in all these constructions depends on the weight of the NTS. Consequently, we improve the weight bound in all these constructions (see Section 4).

An interesting property of our spanners is that they have logarithmic hop-diameter, a result not possible for regular metric spaces. In fact, even 1-dimensional Euclidean space requires linear hop-diameter for light spanners [DES08].

The paper is organized as follows: We close this introductory section with preliminary notes and definitions. We present the first proof in Section 2: We show that there exists a light spanner for all ℓ_p

spaces, and that this implies a light spanner for snowflake metrics as well. The second proof appears in Section 3. Finally, we detail some further applications of our proofs, including TSP algorithms, in Section 4.

1.1 Preliminaries

Let (X, δ) be an arbitrary n-point metric. Without loss of generality, we assume that the minimum inter-point distance of X is equal to 1. We denote by $\Delta = \max_{u,v \in X} \delta(u, v)$ the *diameter* of X.

Doubling dimension. The *doubling dimension* of a metric space (X, δ), denoted by $\mathrm{ddim}(X)$ (or ddim when the context is clear), is the smallest value ρ such that every ball in X can be covered by 2^ρ balls of half the radius. A metric is called *doubling* if its doubling dimension is bounded above by a constant [GKL03].

Hierarchical Nets. A set $Y \subseteq X$ is called an *r-cover* of X if for any point $x \in X$ there is a point $y \in Y$, with $\delta(x, y) \leq r$. A set Y is an *r-packing* if for any pair of distinct points $y, y' \in Y$, it holds that $\delta(y, y') > r$. We say that a set $Y \subseteq X$ is an *r-net* for X if Y is both an r-cover of X and an r-packing. By recursively applying the definition of doubling dimension, we can derive the following key fact [GKL03].

Fact 1.1. *Let $R \geq 2r > 0$ and let $Y \subseteq X$ be an r-packing in a ball of radius R. Then, $|Y| \leq (\frac{R}{r})^{2\,\mathrm{ddim}(X)}$.*

Write $\ell = \lceil \log \Delta \rceil$, and let $\{N_i\}_{i \geq 0}^{\ell}$ be a sequence of hierarchical nets, where $N_0 = X$ and for each $i \in [\ell]$, N_i is a 2^i-net for N_{i-1}. We refer to N_i as the *i-level net*, and the points of N_i are called the *i-level net-points*. Note that $N_0 = X \supseteq N_1 \supseteq \ldots \supseteq N_\ell$, and N_ℓ contains exactly one point. The same point of X may have instances in many nets; specifically, an i-level net-point is necessarily a j-level net-point, for every $j \in [0, i]$. When we wish to refer to a specific instance of a point $p \in X$, which is determined uniquely by some level $i \in [0, \ell]$ (such that $p \in N_i$), we may denote it by the pair (p, i).

Net-tree Spanner. For each $i \in [0, \ell - 1]$, *cross edges* are added between i-level net-points that are within distance $\gamma \cdot 2^i$ from each other, for some parameter $\gamma = \Theta(\frac{1}{\epsilon})$. By Fact 1.1, for each $i \in [0, \ell]$, the degree of any net-point (p, i) in the net-tree spanner (due to i-level cross edges) is $\epsilon^{-O(\mathrm{ddim})}$.

Let R be the sum of radii of all net-points, where the radius $rad(p, i)$ of an i-level net-point (p, i) is equal to 2^i, disregarding the single net-point at level ℓ. It is easy to see that the weight of the net-tree spanner is given by $R \cdot \gamma \cdot \epsilon^{-\Theta(\mathrm{ddim})} = R \cdot \epsilon^{-\Theta(\mathrm{ddim})}$. However, as shown in Section 3.1, R may be as large as $\Theta(\log n) \cdot \omega(MST(X, \delta))$, even for 1-dimensional Euclidean metrics. Consequently, the net-tree spanner in doubling metrics has lightness $\Theta(\log n) \cdot \epsilon^{-\Theta(\mathrm{ddim})}$. On the other hand, it turns out that in *snowflake* doubling metrics – the net-tree spanner is light.

2 Proof via ℓ_p space

In this section we present our first proof that snowflake metrics admit light $(1 + \epsilon)$-spanner. We will first need to prove a certain property concerning Euclidean vectors (Section 2.1). We then show that given a set of d-dimensional vectors S, the edges of a light $(1 + \epsilon)$-spanner for (S, ℓ_2) also form a light $(1 + \epsilon^{O(1)})$-spanner for all (S, ℓ_p) $p \geq 1$, which itself implies a light spanner for snowflake metrics.

Before presenting the proof for ℓ_p spaces, we formally state the Euclidean light spanner theorem of [DHN93]. For an edge set E, let $w_p(E)$ be the sum of the lengths of the edges under ℓ_p. Let $\mathrm{MST}_p(S)$ be the edge set of the minimum spanning tree for S under ℓ_p.

Theorem 2.1. *Let S be a d-dimensional point set of size n. Then there exists an edge set E which forms a $(1 + \epsilon)$-spanner for S under ℓ_2, with $w_2(E) = \epsilon^{-O(d)} w_2(\mathrm{MST}_2(S))$.*

2.1 Properties of Euclidean vectors

In Section 2.2 we will require a property of Euclidean vectors detailed in Lemma 2.3 below. We begin with the following simple fact:

Lemma 2.2. *For $0 \leq \epsilon_0 \leq \epsilon \leq \epsilon_1$, $\epsilon \leq \frac{1}{4}$, and non-negative a, b, if $\epsilon_1 a + \epsilon_0 b \leq \epsilon(a+b)$ then $\epsilon_1 a + \sqrt{\epsilon_0} b \leq \sqrt{\epsilon}(a+b)$.*

Proof: Rewriting both inequalities, we claim that $(\epsilon_1 - \epsilon)a \leq (\epsilon - \epsilon_0)b$ implies that $(\epsilon_1 - \sqrt{\epsilon})a \leq (\sqrt{\epsilon} - \sqrt{\epsilon_0})b$. This claim is confirmed by dividing both sides of the first inequality by $\sqrt{\epsilon} + \sqrt{\epsilon_0}$. The division immediately yields the right hand side of the second inequality; the left hand side follows by noting that $\sqrt{\epsilon} + \sqrt{\epsilon_0} \leq \frac{1}{2} + \frac{1}{2} = 1$, and then observing that $\frac{\epsilon_1 - \epsilon}{\sqrt{\epsilon}+\sqrt{\epsilon_0}} > \frac{\epsilon_1(\sqrt{\epsilon}+\sqrt{\epsilon_0}) - \sqrt{\epsilon}(\sqrt{\epsilon}+\sqrt{\epsilon_0})}{\sqrt{\epsilon}+\sqrt{\epsilon_0}} = \epsilon_1 - \sqrt{\epsilon}$. ∎

Lemma 2.3. *Let V be a set of d-dimensional vectors. Given some vector w, let each vector $v_i \in V$ be decomposed into two orthogonal vectors $v_i^{\perp}, v_i^{\parallel}$, where $v_i^{\perp} + v_i^{\parallel} = v$ and $v_i^{\perp}$ is orthogonal to w. For $0 < \epsilon \leq \frac{1}{4}$, if*

$$\sum_{v \in V} \|v\|_2 = (1+\epsilon) \| \sum_{v \in V} v^{\parallel} \|_2,$$

then

$$\sum_{v \in V} \|v^{\parallel}\|_2 \leq (1+\epsilon) \| \sum_{v \in V} v^{\parallel} \|_2$$

and

$$\sum_{v \in V} \|v^{\perp}\|_2 \leq 3(1+\epsilon)\sqrt{\epsilon} \| \sum_{v \in V} v^{\parallel} \|_2.$$

Proof: The first part of the Lemma is trivial: Since the vectors are orthogonal, $\sum_{v \in V} \|v^{\parallel}\|_2 \leq \sum_{v \in V} \sqrt{\|v^{\parallel}\|_2^2 + \|v^{\perp}\|_2^2} = \sum_{v \in V} \|v\|_2 = (1+\epsilon) \| \sum_{v \in V} v^{\parallel} \|_2$.

Moving to the second part, note that we may assume without loss of generality that all $v_i^{\parallel}$ have the same length: We can always enforce this property by segmenting the v_i's into small vectors with equal parallel contribution without violating the conditions of the lemma.

Define set A, where for each element $a_i \in A$, $a_i = \frac{\|v_i^{\perp}\|_2}{\|v_i^{\parallel}\|_2}$. We have $\|v_i\|_2 = \sqrt{\|v_i^{\parallel}\|_2^2 + \|v_i^{\perp}\|_2^2} = \sqrt{1 + a_i^2} \|v_i^{\parallel}\|_2$. We wish to bound $\sum_{v_i \in V} \|v_i^{\perp}\|_2 = \sum_{v_i \in V} a_i \|v_i^{\parallel}\|_2$. Partition A into two subsets $A_0, A_1 \subset A$, where A_0 contains elements of value less than 1, and A_1 contains elements of value greater or equal to 1. Likewise, partition V into two subsets $V_0, V_1 \subset V$, where $v_i \in V_j$ if and only if $a_i \in A_j$. We have

$$\sum_{v_i \in V_1} \|v_i\|_2 = \sum_{v_i \in V_1} \sqrt{1 + a_i^2} \|v_i^{\parallel}\|_2 > \sum_{v_i \in V_1} (1 + \frac{a_i}{3}) \|v_i^{\parallel}\|_2,$$

$$\sum_{v_i \in V_0} \|v_i\|_2 = \sum_{v_i \in V_0} \sqrt{1 + a_i^2} \|v^{\parallel}\|_2 > \sum_{v_i \in V_0} (1 + \frac{a_i^2}{3}) \|v^{\parallel}\|_2.$$

By the assumption of the lemma,

$$\begin{aligned}
(1+\epsilon) \| \textstyle\sum_{v \in V} v^{\parallel} \|_2 &= \textstyle\sum_{v_i \in V} \|v_i\|_2 \\
&= \textstyle\sum_{v_i \in V_1} \|v_i\|_2 + \sum_{v_i \in V_0} \|v_i\|_2 \\
&> \textstyle\sum_{v_i \in V_1} (1 + \frac{a_i}{3}) \|v_i^{\parallel}\|_2 + \sum_{v_i \in V_0} (1 + \frac{a_i^2}{3}) \|v_i^{\parallel}\|_2 \\
&= \textstyle\sum_{v_i \in V} \|v_i^{\parallel}\|_2 + \sum_{v_i \in V_1} \frac{a_i}{3} \|v_i^{\parallel}\|_2 + \sum_{v_i \in V_0} \frac{a_i^2}{3} \|v_i^{\parallel}\|_2 \\
&\geq \| \textstyle\sum_{v_i \in V} v_i^{\parallel} \|_2 + \sum_{v_i \in V_1} \frac{a_i}{3} \|v_i^{\parallel}\|_2 + \sum_{v_i \in V_0} \frac{a_i^2}{3} \|v_i^{\parallel}\|_2.
\end{aligned}$$

The final inequality follows from the triangle inequality. It follows immediately that $\sum_{v_i \in V_1} \frac{a_i}{3} \|v_i^{\parallel}\|_2 + \sum_{v_i \in V_0} \frac{a_i^2}{3} \|v_i^{\parallel}\|_2 \leq \epsilon \| \sum_{v \in V} v_i^{\parallel} \|_2$. Recall that we may assume that all $\|v_i^{\parallel}\|_2$ are equal, so $\sum_{a_i \in A_1} \frac{a_i}{3} + \sum_{a_i \in A_0} \frac{a_i^2}{3} \leq \epsilon |A| = \epsilon |V|$. Set $\sum_{a_i \in A_1} \frac{a_i}{3} = \epsilon_1 |V_1|$ and $\sum_{a_i \in A_0} \frac{a_i^2}{3} = \epsilon_0 |V_0|$, so that $\epsilon_1 |V_1| + \epsilon_0 |V_0| \leq \epsilon |V|$ and it must be that $\epsilon_1 \geq \frac{1}{3} > \epsilon$ and $\epsilon_0 \leq \epsilon$. If we fix the term ϵ_0 and take the elements of A_0 as variable, we see that $\sum_{a_i \in A_0} a_i$ attains its maximum value when $a_i = \sqrt{3\epsilon_0}$ for all $a_i \in A_0$. So $\sum_{a_i \in A_0} a_i \leq \sqrt{3\epsilon_0} |V_0|$, and using Lemma 2.2 we have that

$$\begin{aligned}
\sum_{v \in V} \|v^{\perp}\|_2 &= \sum_{v_i \in V_1} a_i \|v_i^{\parallel}\|_2 + \sum_{v_i \in V_0} a_i \|v_i^{\parallel}\|_2 \\
&\leq [3\epsilon_1 |V_1| + \sqrt{3\epsilon_0} |V_0|] \|v_0^{\parallel}\|_2 \\
&\leq 3[\epsilon_1 |V_1| + \sqrt{\epsilon_0} |V_0|] \|v_0^{\parallel}\|_2 \\
&\leq 3\sqrt{\epsilon} |V| \|v_0^{\parallel}\|_2 \\
&\leq 3(1+\epsilon)\sqrt{\epsilon} \| \sum_{v_i \in V} v_i^{\parallel} \|_2.
\end{aligned}$$

∎

2.2 Light ℓ_p and snowflake spanners

Here we will show that ℓ_p spaces – and therefore snowflake metrics – admit light spanners. Observe that for any d-dimensional vector x, when $p \geq 2$ we have $\|x\|_p \leq \|x\|_2 \leq d^{\frac{1}{2}-\frac{1}{p}} \|x\|_p$, and when $p \leq 2$ we have $d^{\frac{1}{2}-\frac{1}{p}} \|x\|_p \leq \|x\|_2 \leq \|x\|_p$. We can prove the following lemma:

Lemma 2.4. *Let E be the edge set of a $(1+\epsilon)$-spanner for S under ℓ_2, with weight $w_2(E) = c \cdot w_2(\mathrm{MST}_2(S))$ for some constant $c := c(d, \epsilon)$. Let $d' = \max\{d^{\frac{1}{2}-\frac{1}{p}}, d^{\frac{1}{p}-\frac{1}{2}}\}$. Then set E forms a $(1+\epsilon)(1+3d'\sqrt{\epsilon})$-spanner for S under ℓ_p, $1 \leq p \leq \infty$ with weight $w_p(E) \leq cd' \cdot w_p(\mathrm{MST}_p(S))$.*

Proof: We first prove the weight guarantee. Recall that $\mathrm{MST}_2(S)$ is the minimum weight connected graph under ℓ_2. When $p \geq 2$ we have

$$\begin{aligned}
w_p(E) &\leq w_2(E) \\
&= c \cdot w_2(\mathrm{MST}_2(S)) \\
&\leq c \cdot w_2(\mathrm{MST}_p(S)) \\
&\leq cd^{\frac{1}{2}-\frac{1}{p}} \cdot w_p(\mathrm{MST}_p(S)).
\end{aligned}$$

When $p \leq 2$, we have

$$\begin{aligned}
w_p(E) &\leq d^{\frac{1}{p}-\frac{1}{2}} \cdot w_2(E) \\
&= cd^{\frac{1}{p}-\frac{1}{2}} \cdot w_2(\mathrm{MST}_2(S)) \\
&\leq cd^{\frac{1}{p}-\frac{1}{2}} \cdot w_2(\mathrm{MST}_p(S)) \\
&\leq cd^{\frac{1}{p}-\frac{1}{2}} \cdot w_p(\mathrm{MST}_p(S)).
\end{aligned}$$

This completes the proof of lightness, and we proceed with the stretch guarantee. Consider any vertex pair $x_0, x_t \in S$, connected in E by the minimal (under ℓ_2) edge path $P_{0,t} = \{(x_0, x_1), \ldots, (x_{t-1}, x_t)\}$. Let V be a set of vectors v_i transitioning x_i to x_{i+1}, that is $v_i = x_{i+1} - x_i$. We will employ Lemma 2.3 with respect to V and vector $w = x_t - x_0$. For parallel vectors, we have that $\frac{\sum_{v \in V} \|v^{\parallel}\|_2}{\sum_{v \in V} \|v^{\parallel}\|_p} = \frac{\|\sum_{v \in V} v^{\parallel}\|_2}{\|\sum_{v \in V} v^{\parallel}\|_p}$. Since E is a $(1+\epsilon)$-spanner, Lemma 2.3 gives $\sum_{v \in V} \|v^{\parallel}\|_2 \leq (1+\epsilon) \| \sum_{v \in V} v^{\parallel} \|_2$ and so $\sum_{v \in V} \|v^{\parallel}\|_p \leq (1+\epsilon) \| \sum_{v \in V} v^{\parallel} \|_p$. Lemma 2.3 also gives $\sum_{v \in V} \|v^{\perp}\|_2 \leq 3(1+\epsilon)\sqrt{\epsilon} \| \sum_{v \in V} v^{\parallel} \|_2$: When $p \geq 2$, we have $\sum_{v \in V} \|v^{\perp}\|_p \leq \sum_{v \in V} \|v^{\perp}\|_2 \leq 3(1+\epsilon)\sqrt{\epsilon} \| \sum_{v \in V} v^{\parallel} \|_2 \leq 3d'(1+\epsilon)\sqrt{\epsilon} \| \sum_{v \in V} v^{\parallel} \|_p$, and when $p < 2$ we

have $\sum_{v\in V}\|v^{\perp}\|_p \le d'\sum_{v\in V}\|v^{\perp}\|_2 \le 3d'(1+\epsilon)\sqrt{\epsilon}\|\sum_{v\in V}v^{\parallel}\|_2 \le 3d'(1+\epsilon)\sqrt{\epsilon}\|\sum_{v\in V}v^{\parallel}\|_p$. We conclude that

$$\begin{aligned} w_p(E) &= \sum_{i=0}^{t-1}\|v_i\|_p \\ &\le \sum_{i=0}^{t-1}[\|v_i^{\parallel}\|_p + \|v_i^{\perp}\|_p] \\ &\le (1+\epsilon)\|w\|_p + (3d'(1+\epsilon)\sqrt{\epsilon})\|w\|_p \\ &= (1+\epsilon)(1+3d'\sqrt{\epsilon})\|w\|_p. \end{aligned}$$

∎

It follows that a $(1+\epsilon)$-spanner for (S,ℓ_p) with lightness $\epsilon^{-\tilde{O}(d)}$ can be achieved by building a $(1+O((\epsilon/d')^2))$-spanner for (S,ℓ_2) of lightness $\epsilon^{-O(d)}$, as in Theorem 2.1.

Remark. Lemma 2.4 shows that the stretch guarantee of a Euclidean spanner holds even if the metric is later changed to a different ℓ_p. While this claim is true for Euclidean spanners, it does not hold for all ℓ_p spanners, for example ℓ_∞: Consider three vectors $v_1=(0,0)$, $v_2=(1,1)$, $v_3=(2,0)$, and the two edges $\{v_1,v_2\}$ and $\{v_2,v_3\}$. Then this spanner has no distortion for (S,ℓ_∞), but constant distortion for (S,ℓ_p) and fixed p.

Given a metric $M=(X,\delta)$ with dimension $\mathrm{ddim}(M)$, the metric (X,δ^α) has doubling dimension $\frac{\mathrm{ddim}(M)}{\alpha}$ [GK11]. Har-Peled and Mendel [HPM06] demonstrated that $\frac{1}{2}$-snowflake metrics embed into ℓ_∞ with low distortion and dimension, and their result can be extended to show that (X,δ^α) embeds into d-dimensional ℓ_∞ with distortion $1+\epsilon$ and target dimension $d=\frac{\epsilon^{-O(\mathrm{ddim}(M)/\alpha)}}{1-\alpha}$ [GK11]. Together with Theorem 2.1 and Lemma 2.4, we may conclude:

Theorem 2.5. *Let $M=(X,\delta)$ be an n-point doubling metric with doubling dimension d. For any $0<\alpha<1$ and $\epsilon>0$, there exists a $(1+\epsilon)$-spanner for (X,δ^α) with lightness $2^{\frac{1}{1-\alpha}\epsilon^{-\tilde{O}(\mathrm{ddim}(M)/\alpha)}}$.*

3 Direct snowflake proof

In this section we present a second, tighter proof for the existence of light spanners for snowflake metrics. Let (X,δ) be an arbitrary n-point doubling metric, and let $0<\alpha<1$ be an arbitrary parameter. The corresponding snowflake metric $\tilde{X}:=(X,\delta^\alpha)$ is also doubling, with doubling dimension at most $\frac{\mathrm{ddim}(X)}{\alpha}$. Denote by $\tilde{\Delta}=\max_{u,v\in X}\delta^\alpha(u,v)$ the *diameter* of $\tilde{X}$, and write $\tilde{\ell}=\lceil\log\tilde{\Delta}\rceil$. Next we build the sequence of hierarchical nets $\{\tilde{N}_i\}_{i\ge 0}^{\tilde{\ell}}$ and the net-tree spanner for the snowflake metric $\tilde{X}$; denote by $\tilde{R}$ the sum of radii of all net-points, disregarding the single net-point at level $\tilde{\ell}$.

In what follows we prove that $\tilde{R}=O(\frac{1}{1-\alpha})\cdot\omega(MST(\tilde{X}))$, which implies that the lightness of the net-tree spanner for $\tilde{X}$ is $\epsilon^{-O(\frac{\mathrm{ddim}(X)}{\alpha})}\cdot\frac{1}{1-\alpha}$. We first give some intuition by considering a trivial metric in Section 3.1. The proof of the general case proceeds in two stages.

1. In the first stage (Section 3.2) we construct an auxiliary graph $\tilde{G}$ whose weight $\tilde{W}$ is $\Omega(\tilde{R})$. The graph $\tilde{G}$ will be given as the union of $\tilde{\ell}$ simple paths, and is thus more amenable to analysis than the standard net-tree spanner. The vertex set of this graph $\tilde{G}$ is equal to X.
2. In the second stage (Section 3.3) we show that $\tilde{W}=O(\frac{1}{1-\alpha})\cdot\omega(MST(\tilde{X}))$.

3.1 Intuition and High-Level Ideas

Before delving into the proof, it is instructive to consider the 1-dimensional Euclidean case, and to restrict our attention to $\alpha=\frac{1}{2}$. The intuition behind the general proof is present even in this basic case.

Let $\vartheta=\vartheta_n$ be a set of n points $v_1,\ldots,v_n$ lying on the x-axis with coordinates $1,\ldots,n$, respectively, and consider the corresponding snowflake metric $\tilde{\vartheta}=(\vartheta,\ell_2^{\frac{1}{2}})$. (See Figure 1 for an illustration.) Since the diameter $\tilde{\Delta}$ of $\tilde{\vartheta}$ is $\sqrt{n-1}$, the number of levels in the underlying net-tree is given by $\tilde{\ell}=\lceil\log\tilde{\Delta}\rceil=\lceil\frac{1}{2}\log(n-1)\rceil$. Also, the doubling property implies that the number of i-level net-points is proportional

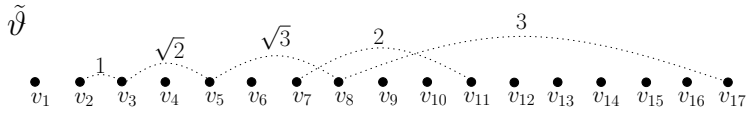

Figure 1: An illustration of the 1-dimensional Euclidean point set $\vartheta = \vartheta_{17} = \{v_1, \ldots, v_{17}\}$, and the corresponding snowflake metric $\tilde{\vartheta} = (\vartheta, \ell_2^{\frac{1}{2}})$. Only 6 edges (out of the total $\binom{17}{2}$ edges) are depicted in the figure, along with their weights which are set according to the snowflake distance function $\|\cdot\|^{\frac{1}{2}}$; thus the weight of edge (v_1, v_{17}), for example, is given by $|17-1|^{\frac{1}{2}} = 4$.

to $\frac{n}{2^{2i}}$. If $\tilde{R}_i$ stands for the sum of radii of all i-level net-points, then $\tilde{R}_i = O(\frac{n}{2^{2i}}) \cdot 2^i = O(\frac{n}{2^i})$, so $\tilde{R} = \sum_{i=0}^{\tilde{\ell}-1} \tilde{R}_i = \sum_{i=0}^{\tilde{\ell}-1} O(\frac{n}{2^i}) = O(n)$. Observe that $\omega(MST(\tilde{\vartheta})) = n-1$. It follows that $\tilde{R} = O(\omega(MST(\tilde{\vartheta})))$, which proves the desired bound on the lightness.[1]

This argument is simple, but it is unclear how to generalize it for arbitrary snowflake metrics. We next provide a more involved proof, whose intuition will be used in the general proof.

Suppose for simplicity of the presentation that $n-1$ is an integer power of 4. For each index $i \in [0, \tilde{\ell}-1]$, we choose a subset $P_i = \{v_1, v_{1+2^{2i}}, v_{1+2\cdot 2^{2i}}, \ldots, v_n\}$ of *pivots* from $\tilde{\vartheta}$, where the ℓ_2 distance between any two consecutive i-level pivots is exactly 2^{2i}. Since $\tilde{N}_i$ is a 2^i-packing, the ℓ_2 distance between any pair of i-level net-points is greater than 2^{2i}. Consequently, it is easy to see that the number $|P_i|$ of i-level pivots is greater than the number $|\tilde{N}_i|$ of i-level and net-points, i.e., $|P_i| > |\tilde{N}_i|$. Let Π_i be the simple path connecting all points of P_i, i.e., $\Pi_i = (v_1, v_{1+2^{2i}}, v_{1+2\cdot 2^{2i}}, \ldots, v_n)$. Notice that the weight of each edge in Π_i (under the snowflake distance function) is equal to the radius 2^i of i-level net-points. Since $|P_i| \geq |\tilde{N}_i|$, the weight $\omega(\Pi_i)$ of Π_i is no smaller than $\tilde{R}_i$. Let $\tilde{G} = \bigcup_{i=0}^{\tilde{\ell}-1} \Pi_i$ be the union of the $\tilde{\ell}$ paths $\Pi_0, \ldots, \Pi_{\tilde{\ell}-1}$, and let $\tilde{W} = \sum_{i=0}^{\tilde{\ell}-1} \omega(\Pi_i)$ be the weight of $\tilde{G}$. Observe that $\tilde{W} \geq \tilde{R}$. (See Figure 2.) We

.

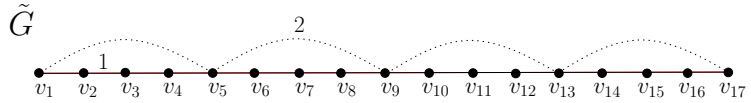

Figure 2: An illustration of $\tilde{G} = \Pi_0 \cup \Pi_1$, in the case that $\vartheta = \{v_1, \ldots, v_{17}\}$ and $\tilde{\ell} = 2$. The path $\Pi_0 = (v_1, v_2, \ldots, v_{17})$ consists of 16 edges of unit weight depicted by solid lines, and the path $\Pi_1 = (v_1, v_5, v_9, v_{13}, v_{17})$ consists of four edges of weight 2 depicted by dotted lines. The weight $\tilde{W}$ of $\tilde{G}$ is equal to $\omega(\Pi_0) + \omega(\Pi_1) = 16 + 8 = 24$.

have thus reduced the problem of lower bounding $\tilde{R}$ to that of lower bounding $\tilde{W}$.

Even though lower bounding $\tilde{W}$ for this basic 1-dimensional case can be done by a direct calculation, our goal here is to present a method that can be applied in the general case.

Any edge (v_l, v_{l+1}) of path $\Pi_0 = (v_1, \ldots, v_n)$ will be called a *path edge*. We say that edge $e = (v_j, v_k)$ (with $1 \leq j < k \leq n$) *loads* a path edge (v_l, v_{l+1}) if $j \leq l < l+1 \leq k$. Thus edge (v_j, v_k) loads the $k-j$ path edges $(v_j, v_{j+1}), \ldots, (v_{k-1}, v_k)$. Next we would like to distribute the weight $\|v_j, v_k\|_2^{\frac{1}{2}} = |k-j|^{\frac{1}{2}}$ of edge $e = (v_j, v_k)$ over all the path edges which it loads. Specifically, the *load* $\xi_{(v_l, v_{l+1})}(e)$ on path edge (v_l, v_{l+1}) caused by edge $e = (v_j, v_k) \in \tilde{G}$, for $j \leq l < l+1 \leq k$, is defined as

$$\xi_{(v_l, v_{l+1})}(e) \;=\; \frac{\|v_j, v_k\|_2^{\frac{1}{2}}}{\|v_j, v_k\|_2} \;=\; \frac{1}{|k-j|^{\frac{1}{2}}}. \tag{1}$$

This means that if we sum up the loads of the $k-j$ path edges $(v_j, v_{j+1}), \ldots, (v_{k-1}, v_k)$ due to edge $e = (v_j, v_k)$, we have the weight $\|v_j, v_k\|_2^{\frac{1}{2}} = |k-j|^{\frac{1}{2}}$ of that edge. For each path edge (v_l, v_{l+1}), the sum of loads on that edge caused by all edges $e \in \tilde{G}$ is called the *load of* (v_l, v_{l+1}) *by* $\tilde{G}$, and is denoted by

[1] This argument does not work for the original (non-snowflake) metric $\vartheta = (\vartheta, \|\cdot\|)$, as there the number of i-level net-points is proportional to $\frac{n}{2^i}$. Thus the sum of radii of all i-level net-points is constant, and therefore $R = \Omega(\log n)$.

$\xi_{(v_l,v_{l+1})} = \xi_{(v_l,v_{l+1})}(\tilde{G})$. The *load of the graph* $\tilde{G}$, $\xi(\tilde{G})$, is the sum of loads over all path edges (v_l, v_{l+1}) by $\tilde{G}$, i.e., $\xi(\tilde{G}) = \sum_{l\in[n-1]} \xi_{(v_l,v_{l+1})}(\tilde{G})$. A double counting argument (see Observation 3.5 below for more details) yields $\tilde{W} = \xi(\tilde{G})$. We have thus reduced the problem of lower bounding $\tilde{R}$ to that of lower bounding the load $\xi(\tilde{G})$ of $\tilde{G}$.

To get some intuition, consider first the graph $\tilde{G}$ from Figure 2, which corresponds to the case $n = 17$. It is easy to see that each path edge is loaded by a single edge of Π_i, for each $i \in [0, 1]$. For example, edge (v_1, v_2) is loaded by edge (v_1, v_2) of path Π_0 and edge (v_1, v_5) of path Π_1, and we thus have

$$\xi_{(v_1,v_2)}(\tilde{G}) \;=\; \xi_{(v_1,v_2)}(v_1, v_2) + \xi_{(v_1,v_2)}(v_1, v_5) \;=\; 1 + \frac{1}{2} \;=\; \frac{3}{2}.$$

Thus $\tilde{W} = \xi(\tilde{G}) = \sum_{l\in[16]} \xi_{(v_l,v_{l+1})}(\tilde{G}) = \frac{3}{2} \cdot 16 = \frac{3}{2} \cdot \omega(MST(\tilde{\vartheta}))$. We turn to the case of general n.

Consider any path edge (v_l, v_{l+1}). It is loaded by a single edge e_i of Π_i, for each level $i \in [0, \tilde{\ell} - 1]$, and its load caused by edge e_i is given by $\xi_{(v_l,v_{l+1})}(e_i) \;=\; \frac{1}{2^i}$. Summing over all $\tilde{\ell}$ levels, we have

$$\xi_{(v_l,v_{l+1})}(\tilde{G}) \;=\; \sum_{i=0}^{\tilde{\ell}-1} \xi_{(v_l,v_{l+1})}(e_i) \;=\; \sum_{i=0}^{\tilde{\ell}-1} \frac{1}{2^i} \;\le\; 2.$$

We conclude that

$$\tilde{W} \;=\; \xi(\tilde{G}) \;=\; \sum_{l\in[n-1]} \xi_{(v_l,v_{l+1})}(\tilde{G}) \;\le\; 2(n-1) \;=\; 2 \cdot \omega(MST(\tilde{\vartheta})).$$

3.2 Stage I

We will now analyze an arbitrary snowflake metric $\tilde{X} = (X, \delta^\alpha)$, where $0 < \alpha < 1$. In this first stage we will construct an auxiliary graph $\tilde{G} = (X, \tilde{E}, \tilde{w})$ whose weight $\tilde{W} = \tilde{\omega}(\tilde{G})$ is at least as large as $\tilde{R}$ (up to a constant). The graph $\tilde{G}$ will be given as the union of $\tilde{\ell}$ simple paths, and is thus more convenient for analysis purposes than the standard net-tree spanner. The vertex set of this graph $\tilde{G}$ is equal to X, and $\tilde{G}$ is equipped with a weight function $\tilde{w}$ which is dominated by the distance function δ^α.

To mimic the 1-dimensional case, we start by computing a Hamiltonian path $\Pi = (v_1, \ldots, v_n)$ for $\tilde{X}$ of weight $\omega(\Pi) = \sum_{i\in[n-1]} \delta^\alpha(v_i, v_{i+1})$ at most $O(\omega(MST(\tilde{X})))$. (Computing such a path Π is a standard procedure, which can be easily carried out given a constant-factor approximation MST for $\tilde{X}$.)

As before, we'd like to compute a set P_i of i-level pivots, for all $i \in [0, \tilde{\ell} - 1]$ – but this process will be carried out more carefully now. Having done that, the graph $\tilde{G}$ will be obtained as before, from the union of the $\tilde{\ell}$ paths $\Pi_0, \ldots, \Pi_{\tilde{\ell}-1}$, with each Π_i connecting the set P_i of i-level pivots via a simple path.

We show how to compute the set P_i of i-level pivots, for $i \in [0, \tilde{\ell}-1]$. Recall that in the 1-dimensional case, any two consecutive i-level pivots are at (snowflake) distance *exactly* 2^i apart; we cannot achieve this property in the general case. Instead, we will make sure that the distance between consecutive pivots will be at least 2^{i-1}. (The reason we use a distance threshold of 2^{i-1} rather than 2^i is technical – this enables us to guarantee that $|P_i| \ge |\tilde{N}_i|$.)

For $i = 0$, we simply take $P_0 = X = \{v_1, \ldots, v_n\}$, and $\Pi_0 = \Pi = (v_1, \ldots, v_n)$. Next consider $i \in [\tilde{\ell}-1]$. The first i-level pivot $p_1^{(i)}$ is v_1. Having assigned the j first i-level pivots $p_1^{(i)}, \ldots, p_j^{(i)}$, the next pivot $p_{j+1}^{(i)}$ is the first point after $p_j^{(i)}$ in Π which is at distance at least 2^{i-1} from it. Formally, let k be the index such that $p_j^{(i)} = v_k$, and let k' be the smallest index after k such that $\delta^\alpha(p_j^{(i)}, v_{k'}) \ge 2^{i-1}$. Then $p_{j+1}^{(i)} = v_{k'}$

For $i \in [0, \tilde{\ell} - 1]$, the i-level path Π_i is a simple path over the i-level pivots. Although an edge e may be very long, we set the weight $\tilde{\omega}(e)$ of each edge e of path Π_i to be 2^{i-1}, and it follows that the edge weights $\tilde{\omega}(e)$ of edges $e \in \Pi_i$ will be dominated by the corresponding snowflake distances. (Indeed, by construction, the snowflake distance between any pair of consecutive i-level pivots is at least 2^{i-1}.) Recall that in the 1-dimensional case we used the weights $\omega(\cdot)$ as given by the snowflake distances – here we use

different weights $\tilde{\omega}(\cdot)$ that are dominated by $\omega(\cdot)$. Denote the edge set of Π_i by $\tilde{E}_i$, with $|\tilde{E}_i| = |P_i| - 1$. Let $\tilde{G} = (X, \tilde{E}, \tilde{\omega})$ be the graph obtained from the union of the $\tilde{\ell}$ paths $\Pi_0, \ldots, \Pi_{\tilde{\ell}-1}$, with $\tilde{E} = \bigcup_{i=0}^{\tilde{\ell}-1} \tilde{E}_i$, and $\tilde{\omega}$ is the weight function as defined above.

As in the 1-dimensional case, we next show that the weight $\tilde{W}$ of $\tilde{G}$ is not much smaller than $\tilde{R}$. The analysis starts with the next observation, which follows immediately from the construction.

Observation 3.1. *Let $p_j^{(i)} = v_k$ and $p_{j+1}^{(i)} = v_{k'}$ be arbitrary consecutive i-level pivots, with $k < k'$. Then for any index $j \in [k, k'-1]$, $\delta^\alpha(p_j^{(i)}, v_j) < 2^{i-1}$. Hence for any $j, \hat{j} \in [k, k'-1]$, $\delta^\alpha(v_j, v_{\hat{j}}) < 2^i$.*

We argue that the number of i-level pivots is no smaller than the number of i-level net-points.

Lemma 3.2. $|P_i| \geq |\tilde{N}_i|$.

Proof: Since the i-level net $\tilde{N}_i$ is a 2^i-packing, any two i-level net-points are at distance at least 2^i apart. By Observation 3.1, for any two consecutive i-level pivots $p_j^{(i)} = v_k$ and $p_{j+1}^{(i)} = v_{k'}$, with $k < k'$, at most one point from $\{v_k, \ldots, v_{k'-1}\}$ belongs to $\tilde{N}_i$. The lemma follows. ∎

Lemma 3.3. *For each $i \in [0, \tilde{\ell} - 1]$, $|\tilde{E}_i| = |P_i| - 1 \geq \frac{1}{2} \cdot |P_i| \geq \frac{1}{2} \cdot |\tilde{N}_i|$.*

Proof: The first equality is immediate from the construction and the third inequality follows from Lemma 3.2. In what follows we prove the second inequality $|P_i| - 1 \geq \frac{1}{2} \cdot |P_i|$, or equivalently $|P_i| \geq 2$.

If all indices $k \in [2, n]$ satisfied $\delta^\alpha(p_1^{(i)} = v_1, v_k) < 2^{i-1}$, the distance between any two points of X would be smaller than $2^i < \tilde{\Delta}$, a contradiction. Let k' be the smallest index for which $\delta^\alpha(p_1^{(i)} = v_1, v_{k'}) \geq 2^{i-1}$. By construction, $v_{k'}$ will be the second i-level pivot $p_2^{(i)}$. Hence $|P_i| \geq 2$, and we are done. ∎

Lemma 3.3 implies that the weight $\tilde{W} = \tilde{\omega}(\tilde{G})$ of $\tilde{G}$ is not much smaller than $\tilde{R}$.

Corollary 3.4. $\tilde{W} \geq \frac{\tilde{R}}{4}$.

Proof: For each $i \in [0, \tilde{\ell} - 1]$, we denote by $\tilde{R}_i$ the sum of radii of all i-level net-points. Observe that $\tilde{R}_i = |\tilde{N}_i| \cdot 2^i$. By Lemma 3.3 and the construction, we have

$$\tilde{W} = \sum_{i \in [0, \tilde{\ell}-1]} \tilde{\omega}(\Pi_i) = \sum_{i \in [0, \tilde{\ell}-1]} |\tilde{E}_i| \cdot 2^{i-1} \geq \sum_{i \in [0, \tilde{\ell}-1]} \frac{1}{2} \cdot |\tilde{N}_i| \cdot 2^{i-1} = \sum_{i \in [0, \tilde{\ell}-1]} \frac{1}{4} \cdot \tilde{R}_i = \frac{\tilde{R}}{4}. \quad ∎$$

3.3 Stage II

In this second stage we show that $\tilde{W} = O(\frac{1}{1-\alpha}) \cdot \omega(MST(\tilde{X}))$. We use a charging scheme, which generalizes the one used in Section 3.1:

1. First, we distribute the weight of each edge of $\tilde{G}$ between the path edges in Π that it "loads". (This is where we leverage on the fact that the metric $\tilde{X} = (X, \delta^\alpha)$ is a snowflake, by using the original distance function δ.)

2. Second, we show that the load incurred in this way by each path edge (v_l, v_{l+1}) is $O(\frac{1}{1-\alpha}) \cdot \delta^\alpha(v_l, v_{l+1})$. This implies that the weight $\tilde{W} = \tilde{\omega}(\tilde{G})$ of $\tilde{G}$ does not exceed the weight $\omega(\Pi)$ of the underlying path Π by more than a factor of $\frac{1}{1-\alpha}$, thereby giving $\tilde{W} = O(\frac{1}{1-\alpha}) \cdot \omega(\Pi) = O(\frac{1}{1-\alpha}) \cdot \omega(MST(\tilde{X}))$.

The *path distance* $\delta_\Pi(v_j, v_k)$ between a pair $v_j, v_k \in X$ of points is given by $\sum_{i=j}^{k-1} \delta(v_i, v_{i+1})$, i.e., the path distance is defined with respect to the original (non-snowflake) distance function δ. As in Section 3.1, any edge (v_l, v_{l+1}) of path $\Pi = (v_1, \ldots, v_n)$ will be called a *path edge*. Also, we say that edge $e = (v_j, v_k)$ (with $1 \leq j < k \leq n$) *loads* a path edge (v_l, v_{l+1}) if $j \leq l < l + 1 \leq k$. Thus edge (v_j, v_k) loads the $k - j$ path edges $(v_j, v_{j+1}), \ldots, (v_{k-1}, v_k)$. Next we would like to distribute the weight $\tilde{\omega}(e)$ of edge $e = (v_j, v_k)$

to all the path edges that it loads. Specifically, the *load* $\xi_{(v_l,v_{l+1})}(e)$ on path edge (v_l, v_{l+1}) caused by edge $e = (v_j, v_k) \in \tilde{G}$, for $j \leq l < l+1 \leq k$, is defined as

$$\xi_{(v_l,v_{l+1})}(e) \;=\; \tilde{\omega}(e) \cdot \frac{\delta(v_l, v_{l+1})}{\delta_\Pi(v_j, v_k)}. \tag{2}$$

(Note that this definition generalizes Equation (1) from Section 3.1.) It is easy to see that the weight $\tilde{\omega}(e)$ of edge e in $\tilde{G}$ is distributed between all the path edges that it loads, so that the sum of loads on these path edges caused by edge e is equal to $\tilde{\omega}(e)$. Also, the load on a specific path edge (v_l, v_{l+1}) caused by edge $e \in \tilde{G}$ is relative to the ratio between the weight of this path edge (with respect to the original metric δ) and the total weight of all the path edges that are loaded by edge e (also with respect to the original metric δ), where the latter term is exactly the path distance between the two endpoints of e.

For each path edge (v_l, v_{l+1}), the sum of loads on that edge caused by all edges $e \in \tilde{G}$, is called the *load of* (v_l, v_{l+1}) *by* $\tilde{G}$, and denoted by $\xi_{(v_l,v_{l+1})} = \xi_{(v_l,v_{l+1})}(\tilde{G})$. The *load of the graph* $\tilde{G}$, $\xi(\tilde{G})$, is the sum of loads over all path edges (v_l, v_{l+1}) by $\tilde{G}$, i.e., $\xi(\tilde{G}) \;=\; \sum_{l\in[n-1]} \xi_{(v_l,v_{l+1})}(\tilde{G})$. A double counting yields:

Observation 3.5. $\tilde{W} \;=\; \sum_{e\in\tilde{G}} \tilde{\omega}(e) \;=\; \sum_{e\in\tilde{G}} \sum_{l\in[n-1]} \xi_{(v_l,v_{l+1})}(e) \;=\; \sum_{l\in[n-1]} \xi_{(v_l,v_{l+1})}(\tilde{G}) \;=\; \xi(\tilde{G})$.

We have thus reduced the problem of lower bounding $\tilde{R}$ to that of lower bounding the load $\xi(\tilde{G})$ of $\tilde{G}$.

We use the following lemma to complete the argument.

Lemma 3.6. *For any index* $l \in [n-1]$, $\xi_{(v_l,v_{l+1})}(\tilde{G}) = O(\frac{1}{1-\alpha}) \cdot \delta^\alpha(v_l, v_{l+1})$.

Proof: Fix any index $l \in [n-1]$. Note that edge (v_l, v_{l+1}) is loaded by a single edge of Π_i, for each $i \in [0, \tilde{\ell}-1]$, denoted e_i. Specifically, edge e_i connects a pair of consecutive i-level pivots $p_j^{(i)} = v_k, p_{j+1}^{(i)} = v_{k'}$, such that $k \leq l < l+1 \leq k'$, and $\delta^\alpha(p_j^{(i)}, p_{j+1}^{(i)}) \geq 2^{i-1}$. Note that all edges of Π_i in $\tilde{G}$ have weight 2^{i-1}, and so $\tilde{\omega}(e_i) = 2^{i-1}$. Thus the load on edge (v_l, v_{l+1}) incurred in level $i \in [0, \tilde{\ell}-1]$ is given by

$$\xi_{(v_l,v_{l+1})}(e_i) \;=\; \tilde{\omega}(e_i) \cdot \frac{\delta(v_l, v_{l+1})}{\delta_\Pi(p_j^{(i)}, p_{j+1}^{(i)})} \;=\; 2^{i-1} \cdot \frac{\delta(v_l, v_{l+1})}{\delta_\Pi(p_j^{(i)}, p_{j+1}^{(i)})},$$

and so the total load on edge (v_l, v_{l+1}) by $\tilde{G}$ is equal to

$$\xi_{(v_l,v_{l+1})}(\tilde{G}) \;=\; \sum_{i\in[0,\tilde{\ell}-1]} \xi_{(v_l,v_{l+1})}(e_i) \;=\; \sum_{i\in[0,\tilde{\ell}-1]} 2^{i-1} \cdot \frac{\delta(v_l, v_{l+1})}{\delta_\Pi(p_j^{(i)}, p_{j+1}^{(i)})}.$$

Define $\eta = \delta^\alpha(v_l, v_{l+1}), t = \lceil \log \eta \rceil$. We first bound the load on edge (v_l, v_{l+1}) incurred in levels $i \in [0, t]$. Since $\delta(v_l, v_{l+1}) \leq \delta_\Pi(p_j^{(i)}, p_{j+1}^{(i)})$, we have $\xi_{(v_l,v_{l+1})}(e_i) \leq 2^{i-1}$. It follows that

$$\sum_{i\in[0,t]} \xi_{(v_l,v_{l+1})}(e_i) \;\leq\; \sum_{i\in[0,t]} 2^{i-1} \;<\; 2\eta \;=\; O(\delta^\alpha(v_l, v_{l+1})).$$

Next, we bound the load on edge (v_l, v_{l+1}) incurred in levels $i \in [t+1, \tilde{\ell}-1]$. We have

$$\delta^\alpha(p_j^{(i)}, p_{j+1}^{(i)}) \;\geq\; 2^{i-1} = 2^t \cdot 2^{i-1-t} \;\geq\; \eta \cdot 2^{i-1-t},$$

and so $\delta(p_j^{(i)}, p_{j+1}^{(i)}) \geq \eta^{\frac{1}{\alpha}} \cdot 2^{\frac{i-1-t}{\alpha}}$. By the triangle inequality, $\delta_\Pi(p_j^{(i)}, p_{j+1}^{(i)}) \geq \delta(p_j^{(i)}, p_{j+1}^{(i)}) \geq \eta^{\frac{1}{\alpha}} \cdot 2^{\frac{i-1-t}{\alpha}}$. Note also that $2^{i-1} = 2^t \cdot 2^{i-1-t} \leq 2\eta \cdot 2^{i-1-t}$. Hence the load on edge (v_l, v_{l+1}) incurred in level $i \in [t+1, \tilde{\ell}-1]$ is given by

$$\xi_{(v_l,v_{l+1})}(e_i) \;=\; 2^{i-1} \cdot \frac{\delta(v_l, v_{l+1})}{\delta_\Pi(p_j^{(i)}, p_{j+1}^{(i)})} \;\leq\; 2\eta \cdot 2^{i-1-t} \cdot \frac{\eta^{\frac{1}{\alpha}}}{\eta^{\frac{1}{\alpha}} \cdot 2^{\frac{i-1-t}{\alpha}}} \;=\; 2\eta \cdot 2^{(i-1-t)(1-\frac{1}{\alpha})}.$$

It follows that

$$\sum_{i\in[t+1,\tilde{\ell}-1]} \xi_{(v_l,v_{l+1})}(e_i) \;\le\; 2\eta\cdot \sum_{i\in[t+1,\tilde{\ell}-1]} 2^{(i-1-t)(1-\frac{1}{\alpha})} \;\le\; \frac{2\eta}{1-2^{(1-\frac{1}{\alpha})}} \;=\; O\left(\frac{1}{1-\alpha}\right)\cdot\delta^{\alpha}(v_l,v_{l+1}).$$

Summarizing, we have

$$\xi_{(v_l,v_{l+1})}(\tilde{G}) \;=\; \sum_{i\in[0,t]} \xi_{(v_l,v_{l+1})}(e_i) + \sum_{i\in[t+1,\tilde{\ell}-1]} \xi_{(v_l,v_{l+1})}(e_i) \;\le\; O\left(\frac{1}{1-\alpha}\right)\cdot\delta^{\alpha}(v_l,v_{l+1}). \qquad \blacksquare$$

Wrapping Up. Recall that $\Pi=(v_1,\ldots,v_n)$ is a Hamiltonian path for $\tilde{X}$ of weight $\omega(\Pi)=\sum_{l\in[n-1]}\delta^{\alpha}(v_l,v_{l+1})$ at most $O(\omega(MST(\tilde{X})))$. Observation 3.5 and Lemma 3.6 imply that

$$\tilde{W} \;=\; \sum_{e\in\tilde{G}}\tilde{\omega}(e) \;=\; \sum_{l\in[n-1]}\xi_{(v_l,v_{l+1})}(\tilde{G}) \;=\; O\left(\frac{1}{1-\alpha}\right)\cdot\sum_{l\in[n-1]}\delta^{\alpha}(v_l,v_{l+1}) \;=\; O\left(\frac{1}{1-\alpha}\right)\cdot\omega(MST(\tilde{X})).$$

By Corollary 3.4, $\tilde{R}\le 4\cdot\tilde{W}=O(\frac{1}{1-\alpha})\cdot\omega(MST(\tilde{X}))$. Finally, recall that the weight of the net-tree spanner for $\tilde{X}$ exceeds $\tilde{R}$ by at most a factor of $\epsilon^{-O(\frac{\mathrm{ddim}(X)}{\alpha})}$, and so its lightness is bounded by $\epsilon^{-O(\frac{\mathrm{ddim}(X)}{\alpha})}\cdot\frac{1}{1-\alpha}$.

Theorem 3.7. *Let $M=(X,\delta)$ be an n-point doubling metric. For any $0<\alpha<1$ and $\epsilon>0$, the lightness of the net-tree spanner for the α-snowflake metric (X,δ^{α}) is $\epsilon^{-O(\frac{\mathrm{ddim}(M)}{\alpha})}\cdot\frac{1}{1-\alpha}$.*

4 Applications

In ICALP'13 Chan *et al.* [CLNS13] showed that, in any doubling metric, one can build a $(1+\epsilon)$-spanner with constant degree and with logarithmic diameter and lightness, within $O(n\log n)$ time. A close examination of the construction of [CLNS13] shows that its lightness is dominated by the lightness of the standard net-tree spanner. Plugging Theorem 3.7 in the construction of [CLNS13] gives rise to a logarithmic improvement in the lightness bound, for snowflake doubling metrics. This result is summarized in the following statement.

Corollary 4.1. *Let $M=(X,\delta)$ be an n-point doubling metric. For any $0<\alpha<1$ and $\epsilon>0$, there exists a $(1+\epsilon)$-spanner for the corresponding α-snowflake metric (X,δ^{α}) with degree $\epsilon^{-O(\frac{\mathrm{ddim}(M)}{\alpha})}$, diameter $O(\log n)$ and lightness $\epsilon^{-O(\frac{\mathrm{ddim}(M)}{\alpha})}\cdot\frac{1}{1-\alpha}$. The runtime of this construction is $\epsilon^{-O(\frac{\mathrm{ddim}(M)}{\alpha})}(n\log n)$.*

Remark. For the 1-dimensional Euclidean metric ϑ_n discussed in Section 3.1, any $(1+\epsilon)$-spanner with diameter $O(\log n)$ must have lightness $\Omega(\log n)$ [DES08]. In contrast, Corollary 4.1 shows that one can get diameter $O(\log n)$ together with constant lightness in snowflake doubling metrics. This reveals a fundamental difference between snowflake and non-snowflake doubling metrics in the context of light spanners.

The construction of [CLNS13] was extended to the fault-tolerant (FT) setting in [CLNS13, Sol13]. Plugging Theorem 3.7 in the FT constructions of [CLNS13, Sol13] gives rise to a logarithmic improvement in the lightness bound, for snowflake doubling metrics. This result is summarized in the following statement.

Corollary 4.2. *Let $M=(X,\delta)$ be an n-point doubling metric. For any $0<\alpha<1$, any $\epsilon>0$ and any integer $0\le k\le n-2$, there exists a k-FT $(1+\epsilon)$-spanner for the corresponding α-snowflake metric (X,δ^{α}) with degree $\epsilon^{-O(\frac{\mathrm{ddim}(M)}{\alpha})}\cdot k$, diameter $O(\log n)$ and lightness $\epsilon^{-O(\frac{\mathrm{ddim}(M)}{\alpha})}\cdot\frac{1}{1-\alpha}(k^2)$. The runtime of this construction is $\epsilon^{-O(\frac{\mathrm{ddim}(M)}{\alpha})}(n\log n+kn)$.*

We turn to algorithms for TSP. A celebrated result of Arora [Aro98] states that for Euclidean TSP, a $(1+\epsilon)$-approximate tour in d-dimensional space can be computed in time $n(\log n)^{\epsilon^{-O(d)}}$, and this result generalizes to ℓ_p spaces as well. Rao and Smith [RS98] utilized a light Euclidean spanner to reduce the runtime to $2^{\epsilon^{-O(d)}}n + 2^{O(d)}n\log n$, and a light ℓ_p spanner is needed to extend their results to ℓ_p space. As a consequence of Lemma 2.4 we have:

Corollary 4.3. *Let X be a set of n d-dimensional vectors. A $(1+\epsilon)$-approximate ℓ_p tour for X can be computed in time $2^{\epsilon^{-\tilde{O}(d)}}n + 2^{\tilde{O}(d)}n\log n$.*

The best runtime for producing a $(1+\epsilon)$-approximate metric tour for an n-point metric X is about $n^{2^{O(\mathrm{ddim}(X))}}$ for fixed ϵ [BGK12]. For the α-snowflake of X, Theorem 3.7 provides a light spanner. Using an appropriate net hierarchy (see [BG13], for example) we can utilize the machinery of Rao and Smith [RS98] and obtain a significantly better result.

Corollary 4.4. *Let $M = (X, \delta)$ be an n-point doubling metric. A $(1+\epsilon)$-approximate metric tour for the α-snowflake of M can be computed in time $2^{\frac{1}{1-\alpha}\epsilon^{-\tilde{O}(\mathrm{ddim}(M)/\alpha)}}n + 2^{\tilde{O}(\mathrm{ddim}(M)/\alpha)}n\log n$.*

Acknowledgments

The authors are grateful to Michael Elkin, Ofer Neiman and Michiel Smid for helpful discussions.

References

[ABN08] Ittai Abraham, Yair Bartal, and Ofer Neiman. Embedding metric spaces in their intrinsic dimension. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '08, pages 363–372. Society for Industrial and Applied Mathematics, 2008.

[ACC+96] Srinivasa Arikati, Danny Z. Chen, L. Paul Chew, Gautam Das, Michiel Smid, and ChristosD. Zaroliagis. Planar spanners and approximate shortest path queries among obstacles in the plane. In Josep Diaz and Maria Serna, editors, *ESA '96*, volume 1136 of *Lecture Notes in Computer Science*, pages 514–528. 1996.

[ADD+93] I. Althöfer, G. Das, D. P. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9:81–100, 1993.

[ADM+95] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. H. M. Smid. Euclidean spanners: short, thin, and lanky. In *Proc. of 27th STOC*, pages 489–498, 1995.

[AMS94] S. Arya, D. M. Mount, and M. H. M. Smid. Randomized and deterministic algorithms for geometric spanners of small diameter. In *Proc. of 35th FOCS*, pages 703–712, 1994.

[Aro98] Sanjeev Arora. Polynomial time approximation schemes for euclidean traveling salesman and other geometric problems. *J. ACM*, 45:753–782, 1998.

[AS97] S. Arya and M. H. M. Smid. Efficient construction of a bounded degree spanner with low weight. *Algorithmica*, 17(1):33–54, 1997.

[Ass83] P. Assouad. Plongements lipschitziens dans $\mathbf{R}^n$. *Bull. Soc. Math. France*, 111(4):429–448, 1983.

[BG13] Yair Bartal and Lee-Ad Gottlieb. A linear time approximation scheme for euclidean tsp. In *FOCS '13*, 2013.

[BGK12] Y. Bartal, L. Gottlieb, and R. Krauthgamer. The traveling salesman problem: low-dimensionality implies a polynomial time approximation scheme. In *Proc. of 44th STOC*, pages 663–672, 2012.

[BRS11] Yair Bartal, Ben Recht, and Leonard J. Schulman. Dimensionality reduction: Beyond the johnson-lindenstrauss bound. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 868–887. SIAM, 2011.

[CG06] H. T.-H. Chan and A. Gupta. Small hop-diameter sparse spanners for doubling metrics. In *Proc. of 17th SODA*, pages 70–78, 2006.

[CLNS13] T.-H. Hubert Chan, Mingfei Li, Li Ning, and Shay Solomon. New doubling spanners: Better and simpler. In *Proc. of the 40th ICALP (1)*, pages 315–327, 2013.

[DES08] Y. Dinitz, M. Elkin, and S. Solomon. Shallow-low-light trees, and tight lower bounds for Euclidean spanners. In *Proc. of 49th FOCS*, pages 519–528, 2008.

[DHM10] Erik D. Demaine, Mohammad Taghi Hajiaghayi, and Bojan Mohar. Approximation algorithms via contraction decomposition. *Combinatorica*, 30(5):533–552, 2010.

[DHN93] G. Das, P. J. Heffernan, and G. Narasimhan. Optimally sparse spanners in 3-dimensional Euclidean space. In *Proc. of 9th SOCG*, pages 53–62, 1993.

[ES13] Michael Elkin and Shay Solomon. Optimal Euclidean spanners: really short, thin and lanky. In *Proc. of the 45th STOC*, pages 646–654, 2013.

[GH12] Michelangelo Grigni and Hao-Hsiang Hung. Light spanners in bounded pathwidth graphs. In Branislav Rovan, Vladimiro Sassone, and Peter Widmayer, editors, *Mathematical Foundations of Computer Science 2012*, volume 7464 of *Lecture Notes in Computer Science*, pages 467–477. Springer Berlin Heidelberg, 2012.

[GK11] Lee-Ad Gottlieb and Robert Krauthgamer. A nonlinear approach to dimension reduction. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 888–899. SIAM, 2011.

[GKL03] A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *Proc. of 44th FOCS*, page 534?543, 2003.

[GR08a] L. Gottlieb and L. Roditty. Improved algorithms for fully dynamic geometric spanners and geometric routing. In *Proc. of 19th SODA*, pages 591–600, 2008.

[GR08b] L. Gottlieb and L. Roditty. An optimal dynamic spanner for doubling metric spaces. In *Proc. of 16th ESA*, pages 478–489, 2008. Another version of this paper is available via `http://cs.nyu.edu/~adi/spanner2.pdf`.

[HPM06] Sariel Har-Peled and Manor Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM J. Comput.*, 35(5):1148–1184, 2006.

[KPX08] Iyad A. Kanj, Ljubomir Perkovi, and Ge Xia. Computing lightweight spanners locally. In *Distributed Computing*, volume 5218 of *Lecture Notes in Computer Science*, pages 365–378. 2008.

[LMN04] James R. Lee, Manor Mendel, and Assaf Naor. Metric structures in l_1: Dimension, snowflakes, and average distortion. In Martn Farach-Colton, editor, *LATIN 2004: Theoretical Informatics*, volume 2976 of *Lecture Notes in Computer Science*, pages 401–412. Springer Berlin Heidelberg, 2004.

[Nei13] Ofer Neiman. Low dimensional embedding of doubling metrics. In *11th Workshop on Approximation and Online Algorithms, September 2013*, 2013.

[NN12] Assaf Naor and Ofer Neiman. Assouads theorem with dimension independent of the snowflaking. *Revista Matematica Iberoamericana*, 28(4):1–21, 2012.

[RS98] Satish B. Rao and Warren D. Smith. Approximating geometrical graphs via "spanners" and "banyans". In *30th annual ACM symposium on Theory of computing*, pages 540–550. ACM, 1998.

[Smi09] M. H. M. Smid. The weak gap property in metric spaces of bounded doubling dimension. In *Proc. of Efficient Algorithms*, pages 275–289, 2009.

[Sol11] S. Solomon. An optimal time construction of Euclidean sparse spanners with tiny diameter. In *Proc. of 22nd SODA*, pages 820–839, 2011.

[Sol13] S. Solomon. From hierarchical partitions to hierarchical covers: Optimal fault-tolerant spanners for doubling metrics. *Technical Report, CoRR abs/1304.8135*, April, 2013.